

EAST Search History

| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|-------|------|---|---|------------------|---------|------------------|
| L1 | 2 | ("20010003839").PN. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:03 |
| L2 | 1134 | (711/141).CCLS. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:03 |
| L3 | 742 | (711/146).CCLS. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:03 |
| L4 | 606 | (711/150).CCLS. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:03 |
| L5 | 456 | (711/148).CCLS. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:04 |
| L6 | 917 | (711/137).CCLS. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:05 |
| L8 | 0 | (tag adj (information file data)) with (" non uniform memory access") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:07 |
| L9 | 48 | (tag\$3 adj (information file data)) and (" non uniform memory access") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:10 |

EAST Search History

| | | | | | | |
|-----|-----|--|---|----|-----|------------------|
| L10 | 184 | (speculat\$6 near4 (read\$3 readout\$3)) with (memor\$5) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:12 |
| L11 | 798 | (" non uniform memory access") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:10 |
| L12 | 4 | 10 and 11 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:12 |
| L13 | 1 | 12 and @ad<"19991209" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:12 |
| L14 | 614 | (speculat\$6 near4 (read\$3 readout\$3)) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:12 |
| L15 | 10 | 14 and 11 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:12 |
| L16 | 6 | 15 and @ad<"19991209" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/03/14 15:19 |


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used **non uniform memory access** **speculative**

 Found **20** of **171,143**

 Sort results by

 Display results
☒ [Save results to a Binder](#)
☒ [Search Tips](#)
☐ Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 20

 Relevance scale ☐ ☐ ☐ ☐ ☐

- 1 [Comparing the memory system performance of the HP V-class and SGI Origin 2000 multiprocessors using microbenchmarks and scientific applications](#)



Ravi Iyer, Nancy M. Amato, Lawrence Rauchwerger, Laxmi Bhuyan

 May 1999 **Proceedings of the 13th international conference on Supercomputing**

Publisher: ACM Press

 Full text available: [pdf\(1.13 MB\)](#)

 Additional Information: [full citation](#), [references](#), [index terms](#)

- 2 [Performance of database workloads on shared-memory systems with out-of-order processors](#)



Parthasarathy Ranganathan, Kourosh Gharachorloo, Sarita V. Adve, Luiz André Barroso

 October 1998 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the eighth international conference on Architectural support for programming languages and operating systems ASPLOS-VIII**, Volume 33 , 32 Issue 11 , 5

Publisher: ACM Press

 Full text available: [pdf\(1.62 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database applications such as online transaction processing (OLTP) and decision support systems (DSS) constitute the largest and fastest-growing segment of the market for multiprocessor servers. However, most current system designs have been optimized to perform well on scientific and engineering workloads. Given the radically different behavior of database workloads (especially OLTP), it is important to re-evaluate key system design decisions in the context of this important class of applicatio ...

- 3 [Architecture and design of AlphaServer GS320](#)



Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

 November 2000 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the ninth international conference on Architectural support for programming languages and operating systems ASPLOS-IX**, Volume 28 , 34 Issue 5 , 5

Publisher: ACM Press


 Full text available: [pdf\(413.91 KB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The

current implementation supports up to 8 such nodes for a total of 32 processors. While
S ...

4 Architecture and design of AlphaServer GS320

 Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren
November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11


Publisher: ACM Press

Full text available:  [pdf\(1.67 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While
S ...

5 Parallel logic programming systems

 Jacques Chassin de Kergommeaux, Philippe Codognet
September 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 3

Publisher: ACM Press


Full text available:  [pdf\(3.51 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Parallelizing logic programming has attracted much interest in the research community, because of the intrinsic OR- and AND-parallelisms of logic programs. One research stream aims at transparent exploitation of parallelism in existing logic programming languages such as Prolog, while the family of concurrent logic languages develops language constructs allowing programmers to express the concurrency—that is, the communication and synchronization between parallel processes—withi ...

Keywords: AND-parallelism, OR-parallelism, Prolog, Warren Abstract Machine, binding arrays, concurrent constraint programming, constraints, guard, hash windows, load balancing, massive parallelism, memory management, multisequential implementation techniques, nondeterminism, scheduling parallel tasks, static analysis

6 Hardware fault containment in scalable shared-memory multiprocessors

 Dan Teodosiu, Joel Baxter, Kinshuk Govil, John Chapin, Mendel Rosenblum, Mark Horowitz
May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(2.05 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Current shared-memory multiprocessors are inherently vulnerable to faults: any significant hardware or system software fault causes the entire system to fail. Unless provisions are made to limit the impact of faults, users will perceive a decrease in reliability when they entrust their applications to larger machines. This paper shows that fault containment techniques can be effectively applied to scalable shared-memory multiprocessors to reduce the reliability problems created by increased mach ...

7 SMTp: An Architecture for Next-generation Scalable Multi-threading

 Mainak Chaudhuri, Mark Heinrich
March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  [pdf\(247.58 KB\)](#) Additional Information: [full citation](#), [abstract](#)


We introduce the SMTp architecture-an SMT processor augmented with a coherence protocol thread context, that together with a standard integrated memory controller can enable the design of (among other possibilities) scalable cache-coherent hardware distributed shared memory (DSM) machines from commodity nodes. We describe the minor changes needed to a conventional out-of-order multi-threaded core to realize SMTp, discussing issues related to both deadlock avoidance and performance. We then compare SMTp p ...

8 The SGI Origin: a ccNUMA highly scalable server

James Laudon, Daniel Lenoski

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  pdf (1.74 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The SGI Origin 2000 is a cache-coherent non-uniform memory access (ccNUMA) multiprocessor designed and manufactured by Silicon Graphics, Inc. The Origin system was designed from the ground up as a multiprocessor capable of scaling to both small and large processor counts without any bandwidth, latency, or cost cliffs. The Origin system consists of up to 512 nodes interconnected by a scalable Craylink network. Each node consists of one or two R10000 processors, up to 4 GB of coherent memory, and ...

9 Parallel and distributed issues

Russell M. Clapp, Trevor Mudge

January 1990 **ACM SIGAda Ada Letters , Proceedings of the working group on Ada performance issues 1990**, Volume X Issue 3

Publisher: ACM Press

Full text available:  pdf (459.54 KB)

Additional Information: [full citation](#), [index terms](#)

10 Efficient shared memory with minimal hardware support

Leonidas I. Kontothanassis, Michael L. Scott

September 1995 **ACM SIGARCH Computer Architecture News**, Volume 23 Issue 4

Publisher: ACM Press

Full text available:  pdf (536.07 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)


Shared memory is widely regarded as a more intuitive model than message passing for the development of parallel programs. A shared memory model can be provided by hardware, software, or some combination of both. One of the most important problems to be solved in shared memory environments is that of cache coherence. Experience indicates, unsurprisingly, that hardware-coherent multiprocessors greatly outperform distributed shared-memory (DSM) emulations on message-passing hardware. Intermediate o ...

11 Space-time scheduling of instruction-level parallelism on a raw machine

Walter Lee, Rajeev Barua, Matthew Frank, Devabhaktuni Srikrishna, Jonathan Babb, Vivek Sarkar, Saman Amarasinghe

October 1998 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the eighth international conference on Architectural support for programming languages and operating systems ASPLOS-VIII**, Volume 33, 32 Issue 11, 5

Publisher: ACM Press

Full text available:  pdf (1.79 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Increasing demand for both greater parallelism and faster clocks dictate that future generation architectures will need to decentralize their resources and eliminate primitives that require single cycle global communication. A Raw microprocessor distributes all of its


resources, including instruction streams, register files, memory ports, and ALUs, over a pipelined two-dimensional mesh interconnect, and exposes them fully to the compiler. Because communication in Raw machines is distributed, com ...

12 Access normalization: loop restructuring for NUMA compilers

 Wei Li, Keshav Pingali

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9


Publisher: ACM Press

Full text available:  [pdf \(1.13 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


In scalable parallel machines, processors can make local memory accesses much faster than they can make remote memory accesses. In addition, when a number of remote accesses must be made, it is usually more efficient to use block transfers of data rather than to use many small messages. To run well on such machines, software must exploit these features. We believe it is too onerous for a programmer to do this by hand, so we have been exploring the use of restructuring compiler technology for ...

13 Systems support for scalable data mining

 William A. Maniatty, Mohammed J. Zaki

December 2000 **ACM SIGKDD Explorations Newsletter**, Volume 2 Issue 2

Publisher: ACM Press

Full text available:  [pdf \(1.13 MB\)](#)

Additional Information: [full citation](#), [index terms](#)


Keywords: KDD, data mining, large data sets, parallelism

14 Implementation machine paradigm for parallel programming

Manohar Rao, Zary Segall, Dalibor Vrsalovic

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  [pdf \(1.05 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Most supercomputers today are parallel computers. In this paper, an approach for efficiently mapping parallel applications onto parallel MIMD machine architectures is introduced. The applicability of this approach to uniform memory access multiprocessors is demonstrated. The paper shows that an intermediate layer of abstraction between the application level and the parallel architecture level is conducive to not only a better software productivity, but also to performance efficient programs. The ...

15 Analytic evaluation of shared-memory systems with ILP processors

 Daniel J. Sorin, Vijay S. Pai, Sarita V. Adve, Mary K. Vernon, David A. Wood


April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available:  [pdf \(1.45 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
[Publisher Site](#)

This paper develops and validates an analytical model for evaluating various types of architectural alternatives for shared-memory systems with processors that aggressively exploit instruction-level parallelism. Compared to simulation, the analytical model is many orders of magnitude faster to solve, yielding highly accurate system performance estimates in seconds. The model input parameters characterize the ability of an application to exploit instruction-level parallelism as well as the interac ...

16 Data distribution support on distributed shared memory multiprocessors
 Rohit Chandra, Ding-Kai Chen, Robert Cox, Dror E. Maydan, Nenad Nedeljkovic, Jennifer M. Anderson

May 1997 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1997 conference on Programming language design and implementation PLDI '97**, Volume 32
Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.55 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Cache-coherent multiprocessors with distributed shared memory are becoming increasingly popular for parallel computing. However, obtaining high performance on these machines requires that an application execute with good data locality. In addition to making effective use of caches, it is often necessary to distribute data structures across the local memories of the processing nodes, thereby reducing the latency of cache misses. We have designed a set of abstractions for performing data distributio ...

17 Comparative performance evaluation of cache-coherent NUMA and COMA architectures
 Per Stenström, Truman Joe, Anoop Gupta


April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture ISCA '92**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.52 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Two interesting variations of large-scale shared-memory machines that have recently emerged are cache-coherent non-uniform-memory-access machines (CC-NUMA) and cache-only memory architectures (COMA). They both have distributed main memory and use directory-based cache coherence. Unlike CC-NUMA, however, COMA machines automatically migrate and replicate data at the main-memory level in cache-line sized chunks. This paper compares the performance of these two classes ...


18 Unified compilation techniques for shared and distributed address space machines
 Chau-Wen Tseng, Jennifer M. Anderson, Saman P. Amarasinghe, Monica S. Lam

July 1995 **Proceedings of the 9th international conference on Supercomputing**

Publisher: ACM Press

Full text available:  [pdf\(1.19 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 Seven-O?lock: A New Distributed GVT Algorithm Using Network Atomic Operations
 David Bauer, Garrett Yaun, Christopher D. Carothers, Murat Yuksel, Shivkumar Kalyanaraman


June 2005 **Proceedings of the 19th Workshop on Principles of Advanced and Distributed Simulation PADS '05**

Publisher: IEEE Computer Society

Full text available:  [pdf\(144.67 KB\)](#)

Additional Information: [full citation](#), [abstract](#)

In this paper we introduce a new concept, network atomic operations (NAOs) to create a zero-cost consistent cut. Using NAOs, we define a wall-clock-time driven GVT algorithm called Seven O?lock that is an extension of Fujimoto's shared memory GVT algorithm. Using this new GVT algorithm, we report good optimistic parallel performance on a cluster of state-of-the-art Itanium-II quad processor systems for both benchmark applications such as PHOLD and real-world applications such as a large-scale T ...

20 Early Experience with Scientific Programs on the Cray MTA-2
 Wendell Anderson, Preston Briggs, C. Stephen Hellberg, Daryl W. Hess, Alexei Khokhlov, Marco Lanzagorta, Robert Rosenberg

November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  [pdf\(233.59 KB\)](#) Additional Information: [full citation](#), [abstract](#)

We describe our experiences porting and tuning three scientific programs to the Cray MTA-2, paying particular attention to the problems posed by I/O. We have measured the performance of each of the programs over many different machine configurations and we report on the scalability of each program. In addition, we compare the performance of the MTA with that of an SGI Origin running all three programs.

Results 1 - 20 of 20

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

+ "non uniform memory access" + "speculative" + "coherency"



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before December 1999

 Terms used **non uniform memory access** **speculative**
coherency

Found 2 of 106,094

Sort results by

relevance

Display results

expanded form

☒ [Save results to a Binder](#)
☒ [Search Tips](#)
☐ Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 2 of 2

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Performance of database workloads on shared-memory systems with out-of-order](#)

[processors](#)

Parthasarathy Ranganathan, Kourosh Gharachorloo, Sarita V. Adve, Luiz André Barroso
 October 1998 **ACM SIGPLAN Notices**, **ACM SIGOPS Operating Systems Review**,
Proceedings of the eighth international conference on Architectural
support for programming languages and operating systems ASPLOS-
VIII, Volume 33, 32 Issue 11, 5

Publisher: ACM Press

Full text available: pdf (1.62 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database applications such as online transaction processing (OLTP) and decision support systems (DSS) constitute the largest and fastest-growing segment of the market for multiprocessor servers. However, most current system designs have been optimized to perform well on scientific and engineering workloads. Given the radically different behavior of database workloads (especially OLTP), it is important to re-evaluate key system design decisions in the context of this important class of applicatio ...

2 [Parallel logic programming systems](#)



Jacques Chassin de Kergommeaux, Philippe Codognet

September 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 3

Publisher: ACM Press

Full text available: pdf (3.51 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Parallelizing logic programming has attracted much interest in the research community, because of the intrinsic OR- and AND-parallelisms of logic programs. One research stream aims at transparent exploitation of parallelism in existing logic programming languages such as Prolog, while the family of concurrent logic languages develops language constructs allowing programmers to express the concurrency—that is, the communication and synchronization between parallel processes—withi ...

Keywords: AND-parallelism, OR-parallelism, Prolog, Warren Abstract Machine, binding arrays, concurrent constraint programming, constraints, guard, hash windows, load balancing, massive parallelism, memory management, multisequential implementation techniques, nondeterminism, scheduling parallel tasks, static analysis

Results 1 - 2 of 2

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

"non uniform memory access" "speculative" "c

[Advanced Search](#)
[Preferences](#)

Web

Results 1 - 10 of about 203 for "**non uniform memory access**" "**speculative**" "**coherency**". (0.39 seconds)

Tip: Save time by hitting the return key instead of clicking on "search"

Non-uniform memory access (NUMA) data processing system that ...

A **non-uniform memory access** (NUMA) computer system includes at least a local ...

Within each processing node 8, status and **coherency** responses are ...

www.freepatentsonline.com/6338122.html - 63k - [Cached](#) - [Similar pages](#)

EP1153349 Internat european software patent - Non-uniform memory ...

Non-uniform memory access (numa) data processing system that speculatively ... [0029]

Within each processing node 8, status and **coherency** responses are ...

gauss.ffii.org/PatentView/EP1153349 - 72k - [Cached](#) - [Similar pages](#)

[PDF] Microsoft PowerPoint - L19 Multiprocessors

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Cache **Coherency**. ¢. A good cache **coherency** protocol can avoid ... **speculative**.
execution. PC. Performance of the Naïve Design ...

www.cs.ucsd.edu/users/calder/classes/win06/240A/slides/L19-Multiprocessors.pdf -

[Similar pages](#)

[PPT] www.hipec.wichita.edu/training/ccNUMA.ppt

File Format: Microsoft Powerpoint - [View as HTML](#)

NUMA: **Non-Uniform Memory Access**. Memory is physically distributed throughout the ...

Cache **Coherency** == System hw guarantees that every cached copy remains ...

[Similar pages](#)

SGI TPL (IRIX 6.5: Developer/OrOn2_PfTune - Glossary)

See **speculative** execution. cache. In common language, a hiding-place for ... Cache
coherency hardware ensures that all cached copies are identical to memory ...

techpubs.sgi.com/library/tp/cgi-bin/getdoc.cgi/

0650/bks/SGI_Developer/books/OrOn2_PfTune/sgi_html/go01.html - 37k -

[Cached](#) - [Similar pages](#)

Citations: The ibm research parallel processor prototype - Pfister ...

60 **Speculative** Lock Elision This chapter presents **Speculative** Lock Elision (SLE) SLE ...

Page Placement For **Non-Uniform Memory Access** Time (NUMA) Shared

citeseer.ist.psu.edu/context/12088/0 - 39k - [Cached](#) - [Similar pages](#)

Citations: Reactive NUMA:A design for unifying S-COMA and CC-NUMA ...

Hardware Support for Extracting Coarse-grain **Speculative**. ... 3.2.5 REACTIVE NON

UNIFORM MEMORY ACCESS (R NUMA) R NUMA dynamically

citeseer.ist.psu.edu/context/344866/0 - 55k - [Cached](#) - [Similar pages](#)

[[More results from citeseer.ist.psu.edu](#)]

Origin2000 Architecture

The Origin uses what is called a directory-based cache **coherency** scheme. ... This
execution will be **speculative** since the condition determining which ...

fy.chalmers.se/UNICC/perf_tuning/O2000Tuning.1.html - 85k - [Cached](#) - [Similar pages](#)

A performance question

bandwidth (**Non-Uniform Memory Access**). And memory access is a prime factor in ... the
shared bus and message passing for cache **coherency**, the speed for Dual ...

einstein.phys.uwm.edu/forum_thread.php?id=1399 - 50k - [Cached](#) - [Similar pages](#)

[PDF] Microsoft PowerPoint - handout04

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Non-Uniform Memory Access (NUMA). – Physically distributed memory. – More
scalable ... **Directory-Based Cache Coherency.** Use indirection ...

www.cs.utexas.edu/users/lin/cs395t/handout04.pdf - [Similar pages](#)

Try your search again on [Google Book Search](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

New! Try out the latest Google Toolbar beta. [Download Now](#) - [About Toolbar](#)




"non uniform memory access" "spec" **Search**

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)



[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

"non uniform memory access" "speculative readout" "coherency"

[Advanced Search](#)
[Preferences](#)

Web

Results 1 - 1 of 1 for "**non uniform memory access**" "**speculative readout**" "**coherency**". (0.69 seconds)

Tip: Try removing quotes from your search to get more results.

[United States Patent Application: 0010003839](#)

[0017] executing the **speculative readout** of the Data from the memories in the ... the processor 10a is configured so that the cache **coherency** is assured by ...

appft1.uspto.gov/.../20010003839&RS=DN/20010003839 - 66k - Supplemental Result -

[Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

New! Try out the latest Google Toolbar beta. [Download Now](#) - [About Toolbar](#)



"non uniform memory access" "speculative readout" "coherency"

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google